Job Polarization, Structural Transformation, and Non-Employment*

Alberto Vindas Q.

Arizona State University

avindasq@asu.edu

March 2019

Abstract

This paper investigates the quantitative importance of the non-employment margin in the labor market outcomes for the United States. During the last 50 years, the labor force has been shifting from producing goods to producing services. In terms of occupations, the routine share decreased, giving way to increases in manual and abstract ones. I argue that these two patterns are related, and that the decrease in non-employment had an important role. I propose and estimate a labor allocation model where goods, market services, and home services use different occupations as inputs. The driving force is productivity growth, which is occupation-specific. Quantitative exercises show that this non-employment channel could significantly slow down polarization and structural transformation, and induce significant displacement within the labor force.

*I am greatly indebted to Bart Hobijn, Gustavo Ventura, and Alex Bick for useful comments, discussions, and support throughout the development of this project. All remaining errors are my own.

1 Introduction

The labor markets in the United States have changed significantly during the last 50 years. The share of routine occupations, heavy on procedural and repetitive tasks, decreased by 28%. On the other hand, manual occupations (heavy on physical tasks and in-person interactions), and abstract occupations (heavy on problem solving and creative tasks) increased by 27 and 45%. In terms of industries, the share of goods fell by 49%, giving way to an increase in services of 33%.[1] The first pattern is related to the *job polarization* process, the shrinking concentration of employment in routine occupations. The second pattern is related to the *structural transformation* process, the reallocation of economic activity across industries.

This transition has received considerable attention due to its implications on wage inequality and mobility costs. Routine occupations tend to be in the middle of the wage distribution, so job polarization pattern translates into higher wage differentials.[2] Human capital has been shown to have a heavy occupation-specific component, so displacements within the labor force would end up in costly adjustments for the workers.[3]

This paper speaks to the latter point. Some of the job polarization analyses frame this transition in terms of job losses and disappearing routine occupations, which ensue large and costly reallocations.[4] Using data from the Current Population Survey, I argue that this is not the case: the adjustment is mostly through decreases in non-employment. Between 1968 and 2018 non-employment decreased from 33 to 22%, a drop of 34%. This translated to an increase in abstract and manual occupations, which explains the decrease in routine's employment share.

This paper also speaks to the occupation-industry mix in the labor force. In quantitative terms, both the occupation mix within industries, and the industry mix within the economy play an important role in explaining overall polarization. Occupational changes within industries have a stronger effect on the increase in

[1] These are percentages with respect to their initial labor shares, so these don't add up to zero.

[2] Goos & Manning (2007) study this pattern in the United Kingdom, Autor & Dorn (2013) focus on the United States, while Goos, Manning & Salomons (2014) expand the analysis for 16 Western European countries.

[3] Kambourov & Manovskii (2009) document this in the United States using data from the Panel Study of Income Dynamics.

[4] Examples include Autor (2010), Acemoglu & Autor (2011), Jaimovich & Siu (2012), and Mandelman (2016).

abstract occupations, while the shift towards services explains most of the increase in manual occupations.

In this article I ask about the quantitative importance of the decrease in non-employment for the productive structure of the economy. To answer that, I propose a labor allocation model explaining the occupational and industrial structure of the economy. It incorporates the non-employment decision, it justifies Baumol's cost disease from an occupational point of view, and gives the polarization process a treatment of the forces taking place between and within broad industries.

This model distinguishes between occupations in labor, and industries in consumption. Its building blocks are motivated by four patterns in the data. First, job polarization has played in a smooth, constant fashion during the last 50 years, so a persistent force should be behind these changes. Second, the adjustment implied a decrease in non-employment, which requires including this margin. Third, the occupational structure within goods and services differs substantially, so the industrial reallocation channel is of quantitative importance. Fourth, both the goods and services industries have polarized similarly, so the forces behind polarization should be occupation, rather than industry-specific.

To study the quantitative implications of the model, I calibrate it to the United States using data from 1968 to 2018. I find that productivity growth is the highest in routine occupations, followed by manual, abstract, and home production. The model is successful in reproducing the occupation dynamics within goods and market services, and is able to generate the movement towards market services we see in the data.

Finally, to assess the quantitative importance of lower non-employment, I perform two counterfactual exercises. The first exercise, inspired by women's insertion into the labor force, freezes non-employment at its 1968 level. This decreases the production of goods and market services by 2 and 18%, holds back structural transformation to its 1999 level, and decreases polarization by an average of 2%. The second exercise, inspired by the home productivity slowdown reported in Bridgman (2016), has home productivity growing at the rate of market services. This increases the production of goods by 17%, decreases the production of market services by 27%, holds back structural transformation to its 1977 level, and decreases polarization by an average of 7.5%.[5] This illustrates not only the importance of this channel, but that its causes also play a important role.

The paper is organized as follows. Section 2 introduces the stylized facts behind job polarization. Section 3 presents the model I use. Section 4 deals with the quantitative matters: first it explains the estimation procedure, and then it goes over the counterfactual exercises. Finally, section 5 concludes.

2 Job Polarization and Non-employment Changes

Between 1968 and 2018, the share of manual occupations in total employment increased by 27 percent, the share of routine occupations decreased by 28 percent, and the share of abstract occupations increased by 45 percent. This is to say, the occupational structure in the United States polarized. These large swings are described as worrisome due to possible displacements to lower paid occupations, as the share of routine occupations decreases, and a higher fraction of the employed population works in lower paying jobs (Jaimovich & Siu, 2012). Meanwhile, non-employment, the fraction of the population that is either unemployed or out of the labor force, decreased 34 percent. These are large changes, and this section documents them with more detail.

Accounting for non-employment's decrease suggests that the net flow of people from non-employment to employment account for polarization, not the flow of people from routine occupations to non-employment. Over this period, the share of the population working in routine occupations remained flat during most of these years, while the share in abstract and manual occupations increased. This is the novel, and main contribution of the empirical analysis: that polarization happened through decreases in non-employment. Few studies focus on this interaction. One notable exception is Cerina, Moro & Rendall (2017), that analyzes how female working hours shaped polarization. Even though the change in non-employment is one of the major adjustment margins, its effect has gone mostly unnoticed in the polarization literature.

I document these changes using data from the Annual Social and Economic (ASEC) supplement to the Current Population Survey. The focus is on the extensive

[5] The comparison points are the model's predicted outcomes for 2018.

margin of labor: whether people work in one of the three groups of occupations, or whether they are non-employed. I consider the population aged between 25 and 65 years, and use their labels to determine industry, occupation (for the employed), and employment status. Lack of hourly data for all these years in CPS precludes studying the hours worked, and therefore, leisure. Fortunately, changes in leisure hours do not point at this being the leading force.[6]

In this paper I follow Cortés et al. (2014) and construct the occupational categories, grouped by their task content. Abstract occupations are intense in tasks that require problem solving, judgment, and creativity. Some examples are managers, lawyers, and architects. Routine occupations are heavy on tasks that follow precise, and well understood routinary procedures. Examples include cashiers, machinists and travel agents. Finally, manual occupations rely more on tasks that require flexibility, in-person interactions, and physical adaptability. These include janitors, bartenders, and nursing aides. Appendix A presents a more detailed discussion of this data source and the classifications.

2.1 Hollowing Out the Employment Distribution

Job polarization has been described as a "hollowing out" of the occupational distribution, since routine occupations are also in the middle of the wage distribution. Indeed, the employment share of routine occupations decreased by 0.33 percentage points per year, which has happened steadily and as early as 1968.

The employment share of routine occupations fell 16.4 percentage points during the 1968-2018 period. Figure 1 shows this by plotting the employment shares for these three occupations. The hollowing out is fairly evident: while manual occupations gained 3 percentage points, abstract occupations gained 13.4. By 2018, the employment shares of abstract and routine occupations were the same.

In addition, these changes have been smooth, and fairly constant over time. A regression analysis summarizes these points more clearly. Table 1 presents the results of regressing the employment shares against a constant term and a trend. All trends

[6] As Cerina, Moro & Rendall (2017) note, most of polarization can be attributed to women, but Aguiar & Hurst (2007), document that leisure has increased almost identically by gender. These studies using time-use surveys also note that total hours of work (including market and home work) are very similar across genders. This suggests that the intensive margin, at the aggregate level, is not the main force driving polarization, and as such is left out of this analysis.

Figure 1: Occupational Job Polarization

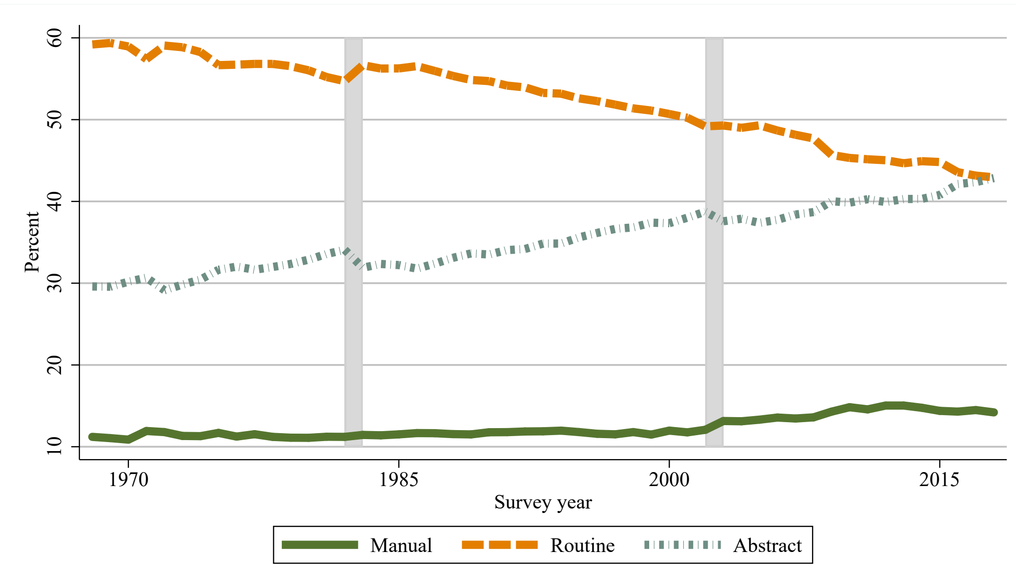

Shaded areas indicate years of major changes in occupational codes. These percentages refer to each occupation's share in employment.
Source: author's calculations using CPS.

Table 1: Employment Share Regression Results

	Constant	Trend	R^2	Std. Error of Estimate
Manual	10.48	0.08	0.74	0.66
Routine	60.51	-0.33	0.96	1.05
Abstract	29.01	0.25	0.96	0.78

These are the results of regressing the occupation's employment share with a constant and a yearly trend term. All coefficients are statistically significant at 1%. The standard error of estimate is the standard deviation of the differences between the observed and predicted shares.
Source: author's calculations using CPS.

are statistically significant, show a good fit to the data with low dispersion around the regression lines. The constant trends mean that job polarization has been happening during the entire period.[7] Several of the studies concerning polarization focus on more recent decades, and posit that it began in the 1980s. Since the CPS data start in 1968, we can go further back in time and state that is has also been happening in earlier years. This is in line with the findings of Bárány & Siegel (2018), who use Census data to document the same pattern for an even longer period.

2.2 Within Industry Job Polarization

Structural transformation, the reallocation of employment across industries, can explain why polarization has been taking place. As industries that are more or less intensive in routine occupations grow, they can drive the occupational employment shares along with them. From an empirical point of view, can trends in this reallocation *between* industries explain most of polarization? In this section, I argue that the answer is no, that changes *within* these industries account for 66% of polarization. This means that most of polarization happens because each industry is demanding less routine workers, and more abstract workers. However, the changes between industries are also quantitatively important, in particular for manual occupations.

The focus of this section are occupational employment shares, but conditional on their industry: the goal is to establish the quantitative importance of the changes within each industry. I organize the productive structure into two industries: goods and services. Following the standard approach, the goods industry consists of agriculture and manufacturing, which include, among other categories, forestry and construction. The services industry includes categories such as retail, and professional and entertainment services.

Graphical examination of within industry occupational changes shows that these are very similar over time. Figure 2 plots the employment shares of each occupation within total employment in each industry. The first panel shows the shares of abstract occupations, which increase over time for both industries. The second shows the opposite effect for routine occupations, and the third shows that manual occupations remain relatively flat, especially in the goods industry. Although the changes were

[7] The regression results for the first-differenced series throw very similar magnitudes, although the coefficient in manual occupations stops being statistically significant.

similar across industries, the magnitudes are larger for services: for the production of goods, the drop in routine occupations was 8.2 percentage points, while it was 11.2 in services. In abstract occupations, the production of goods increased by 8.1 percentage points, and 11.7 in services.

The changes within industries are in line with aggregate changes for routine and abstract occupations. For manual occupations, that is not the case: in the production of goods, its share increased 0.1 percentage points, and decreased 0.5 in the production of services, while overall, its employment share increased by 3. Within industry changes, then, cannot explain the increase in the overall employment share in manual occupations. To account for this, we need to turn to the changes in employment between industries.

Changes between industries can be quantitatively important if the employment shares within each industry are sufficiently different, and if we observe enough reallocation across industries. For manual occupations, both conditions are satisfied. The share of manual occupations in the production of services is 14 times as high as in goods, and the process of structural transformation increased the employment share in services by 19.7 percentage points. This last point is shown in Figure 3, that plots the industry employment shares over time.

To measure the relative importance of these movements, I perform a shift-share decomposition. Overall, changes within industries account for 66% of the occupational shifts, while between industry changes account for the remaining 34%.[8] The changes within industries are most important in increasing abstract and routine occupations' employment shares, while changes between industries matter the most for manual occupations.

The following paragraphs describe this shift-share decomposition. The basic idea comes from expressing the aggregate share of each occupation as a weighted average. For period t:

$$p_t(j) = \sum_I s_t(I) p_t(j|I) \tag{1}$$

where $p_t(j)$ is the economy-wide employment share of occupation j, $s_t(I)$ is the economy-wide employment share of industry I, and $p_t(j|I)$ is the share of occupation

[8] These percentages correspond to the relative contributions of Table 2 weighted by their employment changes.

Figure 2: Occupation Shares Within Industries

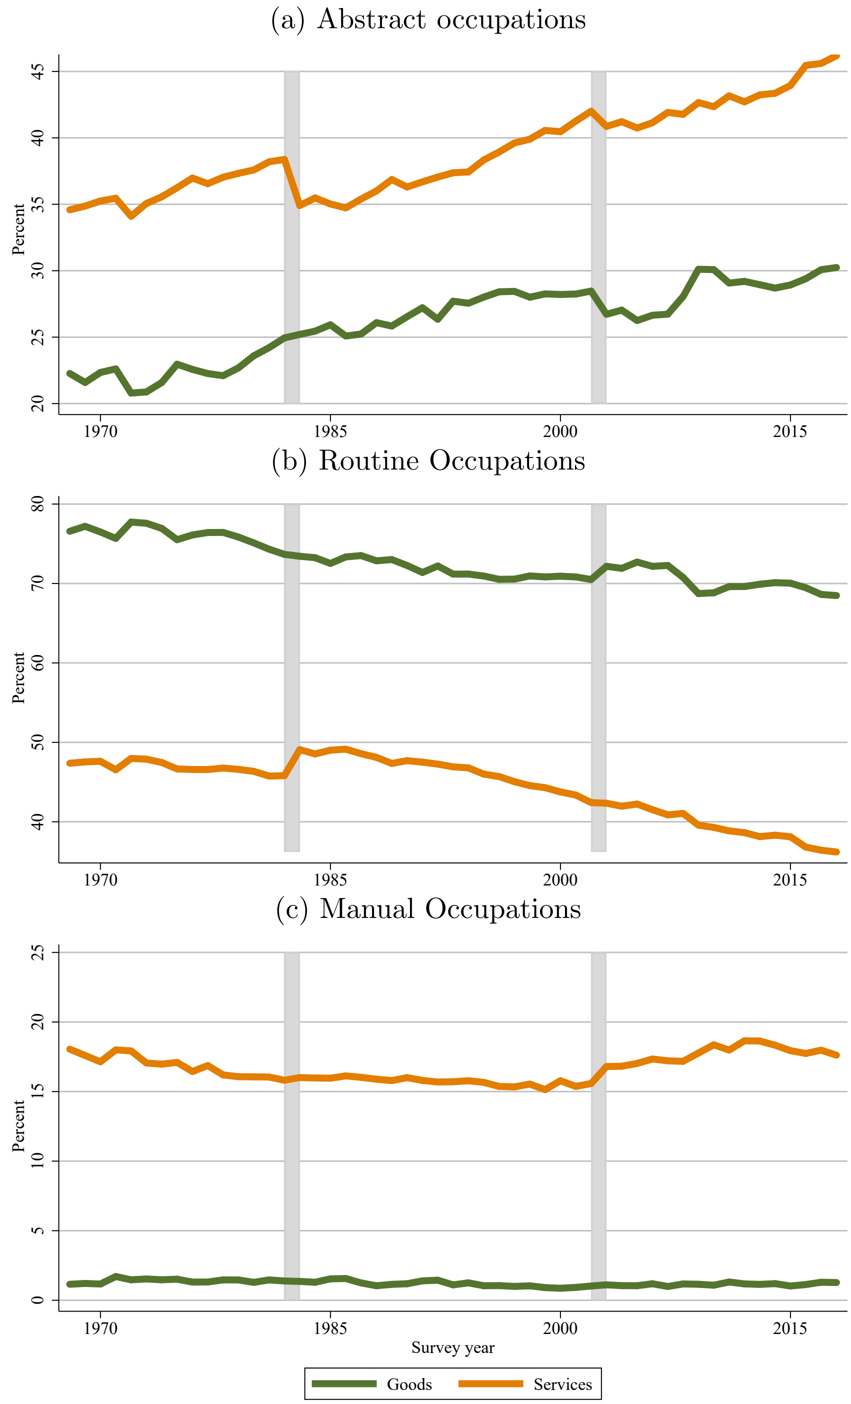

Shaded areas indicate years of major changes in occupational codes. These percentages refer to the share of each occupation in each industry's total labor demand.
Source: author's calculations using CPS.

Figure 3: Industry Shares

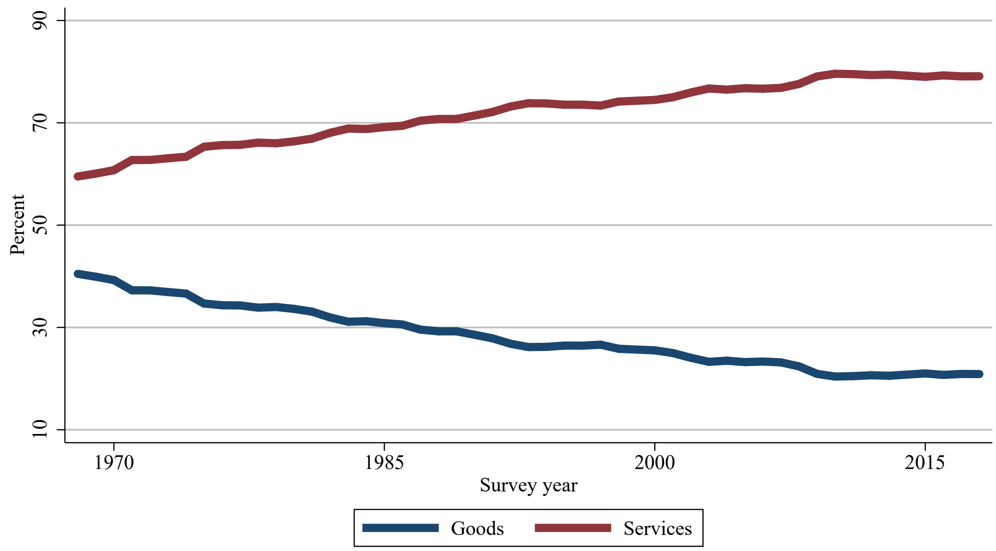

These percentages refer to each industry's share in the labor force.
Source: author's calculations using CPS.

j in industry I.

The change between period 0 and t can be decomposed into a *between* industry effect, and a *within* industry effect. The between industry effect refers to structural transformation. Services is more intensive in abstract and manual occupations than goods, so higher employment shares in services imply higher abstract and manual occupational shares in the economy. The within effect refers to the occupational mix inside each industry. Over time, both the production of goods and services demand more abstract occupations, and less routine ones. In particular:

$$\Delta p_t(j) = \underbrace{\sum_I \Delta s_t(I)\bar{p}(j|I)}_{\substack{\text{Between industries}\\\text{effect}}} + \underbrace{\sum_I \Delta p_t(j|I)\bar{s}(I)}_{\substack{\text{Within industries}\\\text{effect}}} \tag{2}$$

$\bar{p}(j|I)$ is the average between time 0 and t of the conditional occupation share, and $\bar{s}(I)$ that of the industry share. Notice that this decomposition does not have any residual term, since the employment changes are evaluated at their average over the period.

Table 2: Shift-Share Decomposition of Changes in Occupational Shares

Occupation	Absolute Change (p.p.)			Relative Contribution (%)	
	Total Change	Between Industries	Within Industries	Between Industries	Within Industries
Manual	3.8	3.4	0.4	89.5	10.5
Routine	-16.9	-5.7	-11.1	33.7	66.3
Abstract	13.1	2.3	10.8	17.6	82.4

Source: author's calculations using CPS.

Between-industry changes are the most important component for manual occupations, while within-industry changes are the stronger one for routine and abstract occupations. Table 2 shows this. About 90 percent of all of the increase in manual occupations is due to changes between industries, and over 80 percent of the increase in abstract occupations is due to changes within industries. Finally, changes within industries are the stronger component in the decrease of routine occupations. Overall, weighting these relative contributions by the magnitude of the employment share change reveals that 66% of polarization is due to changes within industries, and 34% due to changes between occupations.

In conclusion, the within-industry composition of the economy is the main driver of polarization. It is important to include the industrial composition, however, especially to account for the rise in manual occupations' employment share. This goes in line with previous findings: Autor & Dorn (2013) argue that the expansion of personal services lies behind the increase in low-skill, manual occupations, while Tüzemen & Willis (2013) and Bárány & Siegel (2018) also show that within-industry changes are more important for routine and abstract occupations. These also show similar results when using finer classifications in industries and occupations.

2.3 Lower Non-employment Filled the Edges

The "hollowing out" of the employment distribution can be easily associated to decreases in routine employment, pushing workers into lower paying jobs, and non-employment. In this section, I show that the data is at odds with that interpretation: if we include non-employment as an additional category, routine occupations remain fairly stable as a share of *total population*, as opposed to its decreasing employment share. In addition, the decrease in non-employment allowed for increases in manual and abstract occupations. This means that lower non-employment accounts for job polarization, not decreases in routine employment.

Notice that including non-employed as an additional category changes the reference point: now we study total population aged between 25 and 65 years, rather than employed workers. This allows to study the net flows among these categories, and to better understand the changes leading to polarization.

Figure 4 plots non-employment, and workers in manual, routine, and abstract occupations as shares of total population. There are three main takeaways from this figure: manual and abstract occupations increased their share over time, routine's share remained fairly constant over time, and non-employment decreased for most of this period.

The share of manual and abstract occupations in total population increased during this period, the same way these increased their shares in employment. This means that the increase in the employment shares of manual and abstract occupations responds to a net inflow of workers into these occupations. During these years, the share of these occupations in total population grew by a factor of 1.46 and 1.69, respectively.

The share of routine occupations in employment mostly fell because of lower non-employment: only 26% can be attributed to lower routine employment, and the remaining 74 is due to higher employment in manual and abstract occupations. For this, I use a Taylor approximation of the share in employment of routine occupations $p_t(r)$:

$$p_t(r) = \frac{R_t}{R_t + NR_t} \tag{3}$$

where R_t represents the share in total population of routine occupations, and NR_t represents the share in total population of non-routine occupations. Until 2007, only 9.7% of the decrease in routine occupation's employment share was due to a lower share of routine occupations in total population, and increased to 26% after the great recession of 2008.

Finally, these changes were possible because of the reduction in non-employment. Most of the decreases happened until 2000, and had an increase after the great reces-

Figure 4: Occupational Job Polarization & Labor Non-Participation

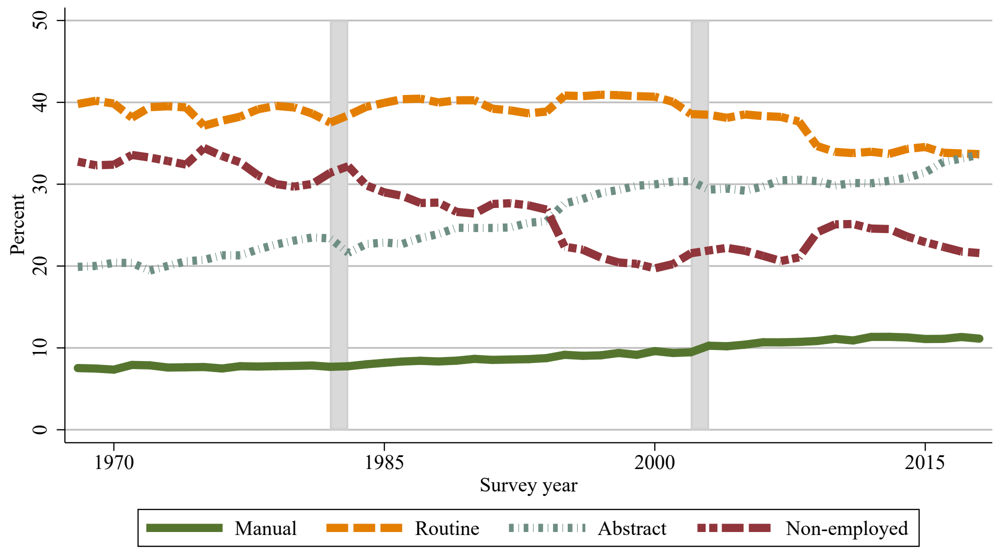

Shaded areas indicate years of major changes in occupational codes. These percentages refer to each category's share in total population.
Source: author's calculations using CPS.

sion that. However, the increases in employment in manual and abstract occupations, that are mostly behind polarization, are the main forces explaining job polarization.

To sum up, in this section I argue that lower non-employment accounts for job polarization, not decreases in routine employment. Net flows out of non-employment, therefore, are key to explain overall polarization. If job displacements were the main cause for the lower employment share in routine occupations, their share in total population would have decreased significantly over time. This is not the case, so we can discard displacements (or increases of it) as the main force behind the fall of routine occupations. This is consistent with the findings in Cortés (2016), where panel data show scant evidence of displacement.

3 A Labor Allocation Model

Section 2 presented a new stylized fact about job polarization: that most of it happened through decreases in non-employment. In this section, I use a model of labor allocation to study quantitatively how changes in non-employment affected it. Its contribution is to provide a general equilibrium framework to analyze the link between non-employment, structural transformation, and job polarization.

Broadly speaking, this is a model of structural transformation where consumers choose between goods and services; goods are produced in the market, and services can be either produced in the market, or at home. In addition, market production demands workers to perform different tasks, which translate into the three occupations studied in the last section. These margins are needed to account for the employment shifts presented in the last section. The driving force is *task specific* technical progress, so the difference in their growth rates induces the three main results: polarization, structural transformation, and changes in non-employment.

The goal of this model is to assess the importance of changes in non-employment for job polarization. The link between these two comes from the consumption of services. Activities that were typically prepared within the household, which correspond to home services in the model, are going through a process of "marketization." This means that they are now being traded in the marketplace, and some of the activities performed as home production are substituted for services produced in the market.[9] This marketization impacts the demand of occupations: higher demand for market services boosts the demand for manual and abstract occupations, while higher demand for home services increases non-employment. This trade-off between market and non-market activities has not gone unnoticed in the literature, but few have focused it towards occupations. The following discussion presents only the relevant aspects of the model; the complete presentation is left to Appendix C.

3.1 Preferences & Technology

This is a discrete-time model where time runs forever. On the consumption side, there are identical households of measure one. These households value three types of consumption: goods that can only be bought in the market, services bought in the market, and services produced at home, as in Ngai & Petrongolo (2014). Differently from them, households only value consumption. The utility level in period t is aggregated according to a nested CES specification:

$$U_t(C_{Gt}, C_{St}) = \left[\omega_G^{\frac{1}{\varepsilon}}(C_{Gt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_S^{\frac{1}{\varepsilon}}(C_{St})^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{4}$$

[9] An early recognition of the importance of home production, and the tendency towards marketization is present in Kuznets (1941).

where C_{Gt} and C_{St} denote the consumption of goods, and a basket of services. Their relative preference weights are ω_G and ω_S, which add up to one, and individually are between zero and one. These two consumption categories consist of broad, and disparate types of consumption, which have an elasticity of substitution of $\varepsilon > 0$.

The basket of services is also represented through a CES aggregator:

$$C_{St} = \left[\varphi_M^{\frac{1}{\eta}} (C_{MSt})^{\frac{\eta-1}{\eta}} + \varphi_H^{\frac{1}{\eta}} (C_{HSt})^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} \tag{5}$$

where C_{MSt} and C_{HSt} denote market and home services. Their relative weights are φ_M and φ_H, which also add up to one, and individually are between zero and one. The types of services provided by these two are much more similar, and their elasticity of substitution is $\eta > 0$.

The production side of this economy has three types of output: goods, market services, and home services. Firms producing goods and market services combine manual, routine, and abstract tasks, while households producing their own services require home workers. Notationally, I denote the sector of production by uppercase letters, so $I \in \{G, MS, HS\}$. Occupations and tasks are denoted by lowercase letters, so $j \in \{m, r, a, h\}$ stand for manual, routine, abstract, and home.

Firms producing goods and market services follow the task approach to production, as in Acemoglu & Autor (2011). Firms, then, combine *tasks* to obtain output, and hire workers to produce those tasks. Ultimately, these are combined through a CES production function:

$$Y_{It} = \left[\sum_{j \in \{m,r,a\}} \alpha_{Ij}^{\frac{1}{\sigma}} (A_{jt} N_{Ijt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{6}$$

where N_{Ijt} denotes the labor input that industry I uses to produce task j, and A_{jt} is the labor productivity in task j. $A_{jt} N_{Ijt}$ then denotes the task input of j in industry I. This task approach follows the same grouping principle for the occupations in the empirical section, therefore workers hired to perform task j are working in occupation j. The elasticity of substitution between the task inputs is $\sigma > 0$, and $\alpha_{Ij} \in (0, 1)$ is the intensity of task j in industry I. Notice that productivity is task-specific, and is the same across industries, as in Duernecker & Herrendorf (2016). Differently to them, this technology is defined over three tasks.

Home services only require home workers to produce, so its technology is represented by a linear function:

$$Y_{HSt} = A_{ht} N_{HSht} \tag{7}$$

where A_{ht} denotes the productivity of home workers, and N_{HSht} denotes the labor input used in home services. These home workers provide the model's counterpart to non-employment.

Labor is homogeneous and perfectly mobile. This means that any two workers can switch between occupations costlessly, and be equally productive if working on the same task. In addition, I normalize the mass of workers to 1, so that all labor inputs in the model match their empirical counterparts as shares in total population.

3.2 Competitive Equilibrium Outcomes

In this section, I analyze the labor reallocation patterns through the lens of the model, assuming perfect competition. Over time, these will only change due to the task-specific productivities. Therefore, I will analyze the productivity changes that must take place to reproduce the employment patterns in Section 2. The model structure requires productivity growth to be the highest in routine tasks, followed by manual, abstract and home production.

This result hinges on three key assumptions related to preferences and technology. In particular, I assume:

1. Tasks to be complements in production.

2. Home and market services to be substitutes in consumption.

3. Goods and services to be complements in consumption.

These are all standard assumptions in the literature, and I will discus these choices in the following subsections. Intuitively, complementarity means that the tasks and categories are very different, and are hard to substitute from one another. The opposite is the case for substitutes: similar categories that are easier to exchange from one another.

Under competitive markets, firms will maximize their profits, and households their utility functions as price takers, as usual. I will analyze the result of these trade-offs around four components: job polarization, industrial productivity, structural transformation, and changes in non-employment.

Explaining Job Polarization

Job polarization within industries results from firms' optimal combination of inputs. Under perfect labor mobility, wages equalize across occupations. Therefore, the demand for workers depends on the productivities of each of the tasks, and their elasticity of substitution.

This is expressed more clearly using firms' first order conditions. In industry I, equalizing the marginal revenue from workers in tasks j and k yields the relative demand of occupations:

$$\frac{N_{Ijt}}{N_{Ikt}} = \frac{\alpha_{Ij}}{\alpha_{Ik}} \left(\frac{A_{kt}}{A_{jt}}\right)^{1-\sigma} \tag{8}$$

This relative occupational demand depends on the relative intensities in that particular industry, and the relative task productivities. The elasticity of substitution, σ, dictates how strongly industries react to changes in productivities.

When occupations are complements, rather than substitutes,[10] firms react to productivity increases by demanding more of the other occupations. Equation (8) shows this more clearly: complementarity implies $\sigma < 1$, and increases in the productivity of task j will boost the relative demand of occupation k. This is due to the cross-productivity effects. If task j increases its productivity, the marginal productivity of the workers in other occupations increases by more, and creates an "excess supply" in occupation j. To benefit the most from this productivity boost, firms demand more of the other occupations.

Recall from Figure 2 that polarization took place within industries. To reproduce that pattern, the growth rate in the productivity of routine tasks has to be the highest

[10] Recall that in this CES formulation, a unitary elasticity of substitution ($\sigma = 1$) implies a Cobb-Douglas production function. As their substitutability decreases, and $\sigma < 1$, inputs are called gross complements. In the limit case, when $\sigma \to 0$, the function converges to a Leontief technology. When $\sigma > 1$, inputs are gross substitutes, given their higher degree of substitutability. When $\sigma \to \infty$, substitutability is perfect and the production function converges to a linear technology.

one, followed by manual, and finally abstract. This is due to the complementarity assumption; it would be reverted if occupations were substitutes.

The productivity terms are modeled in reduced-form way, so these encompass many forces pushing down the demand of routine occupations. Factors affecting these productivities include plain capital accumulation, capital deepening as in Acemoglu & Guerrieri (2008), and increased substitutability with routine workers. The common denominator in these changes is diminishing the cost of performing routine tasks over time. Goos, Manning & Salomons (2014) use a similar approach, but focusing on the occupational cost functions. Their approach on input prices is observationally equivalent to one where the production function is modeled explicitly, as I do in this model.

Treating tasks as complements amounts to assuming that their elasticity of substitution, σ, is less than one. Most of the studies using different labor inputs treat these as substitutes, rather than complements.[11] The task approach to occupations calls for a different interpretation. Since these tasks are very different, it is difficult to substitute for one another. One could think of a production process that combines manual, routine, and abstract tasks, and does so in relatively fixed proportions. This technology is closer to a Leontief specification, which can be represented by a CES production function with a low elasticity of substitution. The effect of increases in productivities, as in Acemoglu & Autor (2011), is to create an "excess supply" of that task, holding the labor composition constant.

Lastly, since productivity changes happen at the task level, it affects all industries equally. This is also observed in the data. In the model, when the elasticity of substitution is equal over industries, the changes in occupational demands should be the same. In the data, comparing the changes in relative occupational shares reveals a similar relationship.[12]

[11] Examples of this are Katz & Murphy (1992), that distinguish between educated and non-educated workers, and Caselli (2015), that distinguishes between experienced and inexperienced workers.

[12] One way to test this relationship is to take the ratio of (8) for a given occupation pair over industries. This shows little variance over time: the coefficients of variation are 0.21 and 0.11, which indicate relative stability.

Explaining Industrial Productivity

The second pattern to focus on is the evolution of labor productivity at the industry level. For the production of goods and market services, occupational demands can be aggregated to a linear technology in total industry demand. At the optimal occupational demands, the production function (18) can be rewritten as:

$$Y_{It} = \tilde{A}_{It} N_{It} \tag{9}$$

where

$$\tilde{A}_{It} = \left[\alpha_{Im} \left(\frac{1}{A_{mt}} \right)^{1-\sigma} + \alpha_{Ir} \left(\frac{1}{A_{rt}} \right)^{1-\sigma} + \alpha_{Ia} \left(\frac{1}{A_{at}} \right)^{1-\sigma} \right]^{\frac{-1}{1-\sigma}} \tag{10}$$

is the average labor productivity in industry $I \in \{G, M\}$, and N_{It} is its associated labor demand.

This productivity is a weighted average of each task's productivity, when labor inputs are combined optimally. With different productivity growth rates in tasks, growth at the industry level will be non-linear, will vary over industries, and will depend on each industry's occupational intensity.

When tasks are complements, and with constant productivity growth, the industry that uses more intensively the task with the highest (lowest) productivity growth increases its overall productivity the most (least). Asymptotically, industry productivity growth rates will converge to the rate of the task with the lowest growth rate. As the share of the occupation with the higher productivity growth decreases, so does its contribution to the growth rate of that industry's productivity. Polarization, then, ends up dampening these productivity gains. At some point in the reallocation process, the share of the occupation with the lowest productivity growth will be so high that its effect on industry productivity will be the only discernible one.

Baumol's cost disease lies at the heart of these dynamics. In their reappraisal of the unbalanced growth model, Baumol, Blackman & Wolff (1985) discuss that the progressivity or stagnancy of economic activities is caused by the technological advances behind their inputs, which correspond to occupations in this setting. The literature in structural transformation has established, by several measures, that goods-producing industries have had higher productivity growth, compared to services-producing industries.[13] This would require goods to be more intensive in routine tasks, and market services to be more intensive in abstract tasks. The data show that this is clearly the case: on average, goods use 63% more routine workers than services, and services use 50% more abstract.

For notational completeness, define the industry equivalent of (10) for home services production:

$$Y_{Ht} = \tilde{A}_{Ht} N_{Ht} \qquad \text{where} \qquad \tilde{A}_{Ht} = A_{ht} \tag{11}$$

Marketization

The third pattern to focus on is the marketization of home production, which relates to the reallocation of productive resources from the home sector to the market.

The decision to consume home services is slightly different than that of goods and market services. To consume home services, households must produce them themselves and give up the market income that they would otherwise earn. The opportunity cost of home production is $p_{Ht} = w_t / \tilde{A}_{Ht}$. In equilibrium, the relative price of home to market services will be inversely related to the sectoral productivities:

$$\frac{p_{Ht}}{p_{Mt}} = \frac{\tilde{A}_{Mt}}{\tilde{A}_{Ht}} \tag{12}$$

Households decide their consumption patterns in services taking into account these relative prices. This translates into the following labor allocations:

$$\frac{N_{Ht}}{N_{Mt}} = \frac{\varphi_H}{\varphi_M} \left(\frac{\tilde{A}_{Mt}}{\tilde{A}_{Ht}} \right)^{1-\eta} \tag{13}$$

Then, again, the evolution over time of this labor allocation will depend on the degree of complementarity and the relative productivity growth. If home and market services are good substitutes, increases in the relative price of home production lead households to substitute its consumption with market services. To do this, they decrease the relative amount of labor dedicated to home production.

This is the marketization result discussed in Freeman & Schettkat (2005): the

[13] A review of this literature, and evidence for several countries is presented in Herrendorf, Rogerson & Valentinyi (2014).

United States has seen a shift of traditional household production to the market. In the model this would reflect a decrease in the relative labor allocated to home production. This, again, calls upon questioning how reasonable the assumptions behind this result are.

Firstly, we should analyze the existence of the home sector itself. Kuznets (1941) was well aware of this fact, and pointed out that incomes within the family economy were a prominent missing item in his estimates of national income. His approximations amounted to more than a quarter of national income in 1929. Hill (1985) uses time-use surveys to establish that for married couples in the mid-1970s, time spent on home work was only slightly behind market work. More recently, Aguiar & Hurst (2007) also find that the amount of hours in home production are substantial, compared to market hours, albeit decreasing over time. These observations should be enough to agree with Benhabib, Rogerson & Wright (1991, p. 1185): "models without home production implicitly make the assumption that the willingness or the incentive of individuals to substitute between market and nonmarket activity is small, but this does not seem to be the conclusion one would want to draw from the evidence."

Secondly, in this model, the time that is not spent in market work is dedicated exclusively to home production. Ngai & Pissarides (2008) posit that in terms of the production carried out in the household, only home services remain. They analyze a much longer time period, and use a structure where home work could be devoted to the three sectors in their study: agriculture, manufacturing, and services. By the late 1920s, they conclude that home production in agriculture and manufacturing was practically gone. Thus, I align with their observations, and assume these productions away.

In this model, the people that are not working in the market are engaging in home production. Time use surveys show that non-employed people engage mostly in home production, performing activities like housework, cooking, and child care (Ngai & Pissarides, 2008). Many of these activities have close counterparts in the market, particularly in the services sector. Other articles building on this idea are Buera & Kaboski (2012) and Ngai & Petrongolo (2014), which I adopt as well.

Thirdly, assuming market and home services are good substitutes in consumption ($\eta > 1$) should not come as a controversial issue. Housework, shopping, food preparation, and caring for other people are among the activities that take most of the time in home production, according to time use surveys. These are all activities that can be easily purchased in the modern marketplace, thus their high degree of substitutability.

Lastly, since market and home services are good substitutes, households will tilt their consumption to the sector with lower price, which is the sector with the higher productivity growth rate. The increase in market participation requires a considerable difference between the growth of home and market productivities. Bridgman (2016) presents evidence to suggest that, effectively, productivity growth in the market has outpaced home productivity, in particular during this paper's period of study.

Structural Transformation

The fourth pattern focuses on the reallocation of consumption and productive resources between market industries, in particular from goods to services. Their relative prices are, again, inversely related to their sectoral productivities:

$$\frac{p_{Gt}}{p_{Mt}} = \frac{\tilde{A}_{Mt}}{\tilde{A}_{Gt}} \tag{14}$$

Preferences are homothetic, so there are no income effects. The expenditure ratio and labor allocation between market goods and services consists of two parts: one that is a price effect, and another one that is a marketization effect. These are represented by:

$$\begin{aligned} \frac{p_{Mt}C_{Mt}}{p_{Gt}C_{Gt}} &= \frac{N_{Mt}}{N_{Gt}} \\ &= \frac{\omega_S}{\omega_G} \underbrace{\left(\frac{\tilde{A}_{Gt}}{\tilde{A}_{Mt}}\right)^{1-\varepsilon}}_{\text{Price effect}} \underbrace{\left\{ \varphi_M^{1-\varepsilon} \left[1 + \frac{\varphi_H}{\varphi_M} \left(\frac{\tilde{A}_{Mt}}{\tilde{A}_{Ht}}\right)^{1-\eta} \right]^{\eta-\varepsilon} \right\}^{\frac{1}{1-\eta}}}_{\text{Marketization effect}} \end{aligned} \tag{15}$$

The price effect responds to relative productivity between *market* industries, while the marketization effect responds to relative productivity between *service* industries.

The price effect behaves similarly to the canonical structural transformation model of Ngai & Pissarides (2007). In this setting, however, task-specific productivity growth is responsible for the growth at the industry level. Complementarity between goods and services (which requires $\varepsilon < 1$) implies that increases in the relative price of services result in increases in its expenditure share.

The marketization effect, on the other hand, is somewhat similar to an income effect. An income effect would induce different consumption patterns, holding constant the productivities across occupations. Within the context of structural transformation, Kongsamut, Rebelo & Xie (2001) introduce income effects with Stone-Geary preferences, and interpret the non-homotheticity term as home production. In this model, as home services become comparatively more expensive, households make up for this by switching out of home production, increasing market work and purchasing more services in the market.

Non-employment

The last component to focus on is the net effect of these forces on non-employment. Structural transformation, with its decrease in the relative price of goods, makes households want to increase their consumption of services. Marketization, on the other hand, makes home production relatively more expensive to market services. Then, the non-employment decision involves a trade-off between home production and market consumption. In equilibrium, the ratio of market employment to non-participation is:

$$\frac{N_{Gt} + N_{Mt}}{N_{Ht}} = \underbrace{\left[1 + \frac{\omega_G}{\omega_S}\left(\frac{\tilde{A}_{St}}{\tilde{A}_{Gt}}\right)^{1-\varepsilon}\right]}_{\text{Structural transformation}} \underbrace{\left[1 + \frac{\varphi_M}{\varphi_H}\left(\frac{\tilde{A}_{Ht}}{\tilde{A}_{Mt}}\right)^{1-\eta}\right]}_{\text{Marketization}} - 1 \tag{16}$$

This expression separates the forces of structural transformation and of marketization. An increase in this ratio implies an increase in labor force participation.

Structural transformation frees up labor that can be used to produce services. This, absent a strong reallocation of consumption within services, would imply higher non-employment. Home production would end up filling part of this increased demand for services. Marketization has the opposing effect, since it becomes relatively cheaper to consume more services from the market.

In the data, there is a sizable decrease in non-employment, most of which is used to fill the increased demand for abstract occupations. This means that the forces of marketization are considerably stronger than those of structural transformation.

Summary of the Model

A brief summary of this model starts with productivity growth rates at the task level, since their differences are the source of the reallocation patterns. The growth rate is higher in routine tasks, followed by manual, abstract, and finally home production. The interaction between these productivities, technology, and preferences yields five outcomes:

Polarization: firms demand more workers in the abstract and manual occupations, relative to routine, because tasks are complements in production.

Industry Productivity Growth: productivity growth is higher in goods because it is more intensive in routine tasks than market services.

Marketization: households work more in the market and substitute home with market services because of increasing opportunity costs of home production.

Structural Transformation: households demand more services because the relative price of goods decreases.

Non-employment: non-employment decreases because the effect of marketization dominates the effect of structural transformation on services.

4 Quantitative Results: How Non-employment Affects Polarization

This section explores the quantitative side of the model. First, I describe the estimation procedure, and analyze its results. With the estimation in hand, I conduct two counterfactual exercises to assess the importance of the increase in labor force participation.

4.1 Calibrating the Model

In this section, I explain briefly how to calibrate the model, and the moments I use. To begin with, I choose the two elasticities in the utility function based on previous studies. Based on Duernecker & Herrendorf (2016), I set $\varepsilon = 0.05$ (between goods and compound services). Based on Rogerson (2009) and Ngai & Petrongolo (2014), I set $\eta = 2.3$ (between market and home services).[14]

With these restrictions, there are eleven time-invariant parameters and four terminal conditions to determine: six task intensities (three for each market industry), the labor elasticity of substitution in market production, the four final task productivities, and the four preference weights in the consumption (for goods and services, and for home and market services). Assuming constant growth rates, there is only need to look at the initial and final years, which are denoted here by $t = 0$ and $t = T$.[15] I back out their estimates from US data following these steps:

1. Impose the normalization $A_{m0} = A_{r0} = A_{a0} = A_{h0} = 1$.
2. Use the initial market occupation shares N_{Ij0} to solve for α_{Ij}.
3. Use the initial home and market services shares N_{H0}, N_{M0} to solve for φ_H and φ_M.
4. Use the initial market goods and services shares N_{G0}, N_{M0} to solve for ω_G and ω_S.
5. Use the final employment shares in the market services industry N_{MjT} to solve for final relative productivities $(A_{aT}/A_{rT})^{1-\sigma}$ and $(A_{mT}/A_{rT})^{1-\sigma}$.
6. Use the final relative employment share N_{MT}/N_{ST} to solve for the labor elasticity of substitution σ.
7. Use the growth factor of real per capita GDP to solve for A_{rT}.
8. Use the final home and market services shares to solve for A_{hT}

Further details of this procedure are discussed in Appendix D. Table 3 shows the time-invariant parameters of the model, and Table 4 the occupation and industry productivity estimates.

[14] This elasticity is in the high end of the estimates available. It was purposely chosen as such, because these come from studies looking at the substitution between home and *total* market goods. This selection attempts to make up for goods being included.

[15] Notice this same procedure could be used to infer a more detailed productivity path on a yearly basis, but the smooth labor share paths suggest this is a reasonable assumption.

Table 3: Model Parameters

	Parameter	Value	Source
ε:	Elasticity of substitution in utility b/w goods and combined services	0.05	Herrendorf, Rogerson & Valentinyi (2013)
η :	Elasticity of substitution in utility b/w home and market services	2.30	Ngai & Petrongolo (2014)
ω_G:	Preference weight for goods in utility	0.27	Initial goods and services relative price
ω_S:	Preference weight for combined services in utility	0.73	Initial goods and services relative price
φ_H:	Preference weight for home services in utility	0.45	Initial home and market services labor shares
φ_M:	Preference weight for market services in utility	0.55	Initial home and market services labor shares
σ :	Labor elasticity of substitution in production b/w market occupations	0.22	Final employment shares in market services
α_{Ga}:	Intensity of abstract occupations in market goods production	0.22	Initial industry-specific occupation shares
α_{Gr}:	Intensity of routine occupations in market goods production	0.77	Initial industry-specific occupation shares
α_{Gm}:	Intensity of manual occupations in market goods production	0.01	Initial industry-specific occupation shares
α_{Ma}:	Intensity of abstract occupations in market services production	0.35	Initial industry-specific occupation shares
α_{Mr}:	Intensity of routine occupations in market services production	0.47	Initial industry-specific occupation shares
α_{Mm}:	Intensity of manual occupations in market services production	0.18	Initial industry-specific occupation shares

All parameters, except the first two elasticities, are estimated from CPS data. See section 4.1 and appendix D for more details.

Table 4: Productivity Estimates

		Year		Average
		1968	2018	Growth Rate
A_h:	Home production	1	1.01	0.03%
A_a:	Abstract occupations	1	1.32	0.56%
A_r:	Routine occupations	1	2.59	1.96%
A_m:	Manual occupations	1	1.93	1.35%
$\tilde{A}_G$:	Market goods industry	1	2.15	1.57%
$\tilde{A}_M$:	Market services industry	1	1.87	1.29%

These results are in line with those explained in the discussion section, so the qualitative predictions remain. Now we can comment on their quantitative side. As expected, growth in all occupation-specific productivities is positive. Productivity growth is such that by 2018, a worker in routine occupations is 35% more productive than a worker in manual occupations, and almost twice as much than a worker in abstract occupations. This goes in line with the routinization hypothesis, but established as a force working since (at least) the beginning of the sample. The elasticity of substitution between occupations is considerably lower than other estimates. Again, this stems from considering occupations as different factors of production, which is induced by the task-oriented grouping.

Table 5 shows the model's predictions with the estimated parameters. By design, it is able to match all of the labor shares in 1968, and the relative labor shares within services (for the three occupations in market services, plus the ratio between market and home services). The model is fairly successful at reproducing the relative occupation shares in the goods industry, and a little less so in reproducing the decrease in the goods share. The normalized root-mean-squared-error (RMSE) shows that for a model with constant growth rates, it is still satisfactory in capturing the main forces driving the changes in labor markets.

4.2 Assessing the Importance of Non-employment

One of the main goals of this paper is to assess the importance of lower non-employment in the distribution of employment in the United States. The model at hand allows us to do this, and study its effects on structural transformation, on

Table 5: Actual and Fitted Labor Shares

	1968		2018		Normalized
	Data	Model	Data	Model	RMSE
Share of total population in:					
Goods production	27.3	27.3	16.3	20.9	15.0%
Market services production	40.1	40.1	60.7	58.7	7.5%
Home services production	32.6	32.6	23.0	20.4	8.7%
Share of goods labor demand in:					
Manual occupations	1.2	1.2	1.0	1.3	18.9%
Routine occupations	76.7	76.7	70.5	65.6	1.9%
Abstract occupations	22.1	22.1	28.5	33.2	8.1%
Share of market services labor demand in:					
Manual occupations	18.1	18.1	17.7	17.6	8.7%
Routine occupations	47.4	47.4	38.3	36.2	8.0%
Abstract occupations	34.5	34.5	43.9	46.2	5.3%

Normalized root-mean-squared-error is calculated using the average of the time series.

overall polarization, and on sectoral output. To do that, I study two experiments. In the first one, I restrict the model so that non-employment stays at its 1968 level. In the second one, I increase the productivity growth in home production to shut down the marketization channel. The comparison point of these exercises are the output levels and labor allocations of the baseline model in 2018. The first year of my sample coincides with the years when labor force participation started to increase. Juhn & Potter (2006) report that between 1948 and 1968 the participation rate remained relatively stable, and after that it increased. This provides a convenient turning point to perform counterfactual exercises.

Freezing Non-Employment

The expansion of the labor force, which has lowered non-employment, has drawn a considerable amount of attention recently, and has been associated with the insertion of women into the workplace. Several interpretations have been presented to explain this. One of them focuses on the social attitudes towards women's work: Fernández, Fogli & Olivetti (2004) discuss the gradual transformation of the family model (where working mothers set an example for future generations), while Goldin (2006) discusses several other social changes fueling the "quiet revolution" that made women think in terms of lifelong careers. In the spirit of these articles, I freeze non-employment at its level in 1968. One interpretation of this counterfactual would be to consider how different the productive structure would have been, had societal attitudes towards women's work not changed as much. The results of this are in Table 6.

Table 6: Freezing Non-employment, Counterfactual Results

	Baseline model	Counterfactual prediction	Change (%)
Non-employment	0.20	0.33	59.4
Goods output	0.46	0.45	-1.8
Market services output	1.08	0.87	-19.2
Share in labor force of:			
Goods industry	0.26	0.30	15.5
Services industry	0.74	0.70	-5.5
Share in total population of:			
Manual occupations	0.13	0.13	-5.0
Routine occupations	0.44	0.45	2.7
Abstract occupations	0.43	0.42	-1.2

This are the model's predictions for 2018. The counterfactual exercise freezes the labor force participation rate at its 1968 level, without assuming different rates of productivity growth. Percent changes are reported with respect to the baseline model's predictions.

Holding non-employment fixed decreases output in both market sectors, disproportionately so for market services. This is not surprising, since this experiment deliberately holds down the inputs for market production, and restricts households to consume more home services than they would otherwise want. Within the workforce, this slows down structural transformation: the employment share of goods ends up being 16% higher, despite market productivities remaining the same. This division of labor between goods and services makes the productive structure look like the model's prediction for 1999. Focusing on job polarization, the forces acting *within* both industries would still take place, so the labor force would polarize to a similar degree. The main difference is that in this case, the entirety of the adjustment takes place *within* the labor force, instead of through an expansion of it. Then, worker displacement would be a significant source of adjustment.

Increasing Home Productivity

Another explanation for the change in non-employment focuses on the productivity of the household sector. In this vein, the second experiment keeps non-employment from decreasing, but because of very different reasons. Bridgman (2016) suggests that productivity growth in the home sector had a considerable slowdown, coinciding with the initial years of this CPS sample. He also reports that previous to these years, home productivity had been growing at a similar pace to market activities. This is consistent with non-employment being stable before 1968, and increasing after this break. This begs the question: if the slowdown in home productivity was the sole cause of marketization, what would have happened without it? That is, how different would the productive structure look like if productivity in the home production sector had grown at the same rate as market services? The outcome of this experiment is in Table 7.

Table 7: Increasing Home Productivity, Counterfactual Results

	Baseline model	Counterfactual prediction	Change (%)
Non-employment	0.20	0.34	65.3
Goods output	0.47	0.56	18.2
Market services output	1.14	0.81	-29.2
Share in labor force of:			
Goods industry	0.26	0.37	42.0
Services industry	0.74	0.63	-14.9
Share in total population of:			
Manual occupations	0.13	0.12	-13.5
Routine occupations	0.44	0.47	-7.4
Abstract occupations	0.43	0.41	-3.4

This are the model's predictions for 2018. The counterfactual exercise shuts down the marketization channel by increasing home production's productivity at the same pace as market services. Percent changes are reported with respect to the baseline model's predictions.

The results in terms of output are more dramatic in this case: production in goods increases, while in market services it decreases even more. Households are much more efficient in producing their own services, so they do not increase their reliance on the market. This explains the greater hit that market services take.

Non-employment increases; since I am purposely shutting down the marketization channel, only structural transformation has an effect on the consumption of total services (equation (16) shows this). These effects add up to goods having an even higher share in market employment, which increases by almost 40%. The resulting division of labor makes the productive structure look like the model's prediction for 1977. In terms of polarization, the intra-industry reallocation would still take place. In this case, however, lower growth in market services leads to less polarization than in the first counterfactual. As in the first counterfactual, all the reallocation would happen within the labor force.

These two exercises illustrate how important the decrease in non-employment is for both structural transformation and job polarization. In the two cases, both structural transformation and polarization are slowed down by the induced dynamics of home production. This raises the question of, for instance, to what extent the differences in the labor market structure of the United States and some European countries are driven by the channels highlighted in this model. Prescott (2004) explores the issue of differences in hours worked by looking at taxes. Home production could play a significant role in this setting. If cultural preferences favor home production, despite a lower productivity in it, it could look more like the first counterfactual exercise. If, on the other hand, European households did not experience a significant decrease in their home productivity growth, it could look more like the second counterfactual. These two cases show very different implications for growth and the causes for the differences in their productive structures. Polarization slows down, but the road it takes is quite different. Unfortunately, this model is unable to speak of mobility costs by design. Displacement costs would play an important role in terms of welfare, and could possibly slow down the adjustment process. This interesting avenue of research is left for other research projects.

5 Conclusions

In this article I study three recent trends in the United States labor market: job polarization, structural transformation, and decreases in non-employment. With the goal of quantifying the importance of the non-participation margin, I propose a labor allocation model to explain this occupational and industrial structure.

Quantitatively, the model implies higher growth in the occupation of routine occupations, followed by manual, abstract, and finally home production. It is able to reproduce the occupational structure within industries, and the shift towards market services.

Counterfactual exercises suggest that this expansion is very important. Holding constant the attitudes that allowed lower non-employment decreases the production of goods and market services by 2 and 19%, holds back structural transformation to its 1999 level, and decreases polarization by an average of 3%. The second exercise, inspired by the home productivity slowdown reported in Bridgman (2016), has home productivity growing at the rate of market services. This increases the production of goods by 18%, decreases the production of market services by 29%, holds back structural transformation to its 1977 level, and decreases polarization by an average of 8%.

Future work includes extensions to this model to incorporate the age structure of the economy through an overlapping generations setting, considering the gender component to explain the wide differences in their labor market outcomes, and inducing labor heterogeneity to include an analysis of wage polarization.

References

Acemoglu, Daron, & David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." , ed. O. Ashenfelter & D. Card Vol. 4 of *Handbook of Labor Economics*, Chapter 12, 1043–1171. Elsevier.

Acemoglu, Daron, & Veronica Guerrieri. 2008. "Capital Deepening and Non-balanced Economic Growth." *Journal of Political Economy*, 116(3): 467–498.

Aguiar, Mark, & Erik Hurst. 2007. "Measuring Trends in Leisure: The Allocation of Time over Five Decades." *The Quarterly Journal of Economics*, 122(3): 969–1006.

Autor, David. 2010. "The Polarization of Job Opportunities in the U.S. Labor Market: Implications for Employment and Earnings." The Hamilton Project and Center for American Progress.

Autor, David, & David Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market." *American Economic Review*, 103(5): 1553–97.

Baumol, William, Sue Anne Batey Blackman, & Edward N. Wolff. 1985. "Unbalanced Growth Revisited: Asymptotic Stagnancy and New Evidence." *The American Economic Review*, 75(4): 806–817.

Benhabib, Jess, Richard Rogerson, & Randall Wright. 1991. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *Journal of Political Economy*, 99(6): 1166–1187.

Bridgman, Benjamin. 2016. "Home Productivity." *Journal of Economic Dynamics and Control*, –.

Buera, Francisco, & Joseph Kaboski. 2012. "The Rise of the Service Economy." *The American Economic Review*, 102(6): 2540–2569.

Bárány, Zsófia L., & Christian Siegel. 2018. "Job Polarization and Structural Change." *American Economic Journal: Macroeconomics*, 10(1): 57–89.

Caselli, Francesco. 2015. "Experience-Biased Technical Change." Centre for Macroeconomics (CFM) Discussion Papers 1517.

Cerina, Fabio, Alessio Moro, & Michelle Petersen Rendall. 2017. "The Role of Gender in Employment Polarization." Centre for Macroeconomics (CFM) Discussion Papers 1704.

Cortés, Guido Matias, Nir Jaimovich, Christopher Nekarda, & Henry Siu. 2014. "The Micro and Macro of Disappearing Routine Jobs: A Flows Approach." National Bureau of Economic Research, Inc NBER Working Papers 20307.

Cortés, Guido Matías. 2016. "Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data." *Journal of Labor Economics*, 34(1): 63 – 105.

Duernecker, Georg, & Berthold Herrendorf. 2016. "Structural Transformation of Occupation Employment." University of Mannheim & Arizona State University Working Paper.

Fernández, Raquel, Alessandra Fogli, & Claudia Olivetti. 2004. "Mothers and Sons: Preference Formation and Female Labor Force Dynamics." *The Quarterly Journal of Economics*, 119(4): 1249–1299.

Freeman, Richard, & Ronald Schettkat. 2005. "Marketization of Household Production and the EU-US Gap in Work." *Economic Policy*, 20(41): 7–50.

Goldin, Claudia. 2006. "The Quiet Revolution That Transformed Women's Employment, Education, and Family." *American Economic Review*, 96(2): 1–21.

Goos, Maarten, & Alan Manning. 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *The Review of Economics and Statistics*, 89(1): 118–133.

Goos, Maarten, Alan Manning, & Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review*, 104(8): 2509–26.

Herrendorf, Berthold, Richard Rogerson, & Ákos Valentinyi. 2013. "Two Perspectives on Preferences and Structural Transformation." *American Economic Review*, 103(7): 2752–89.

Herrendorf, Berthold, Richard Rogerson, & Ákos Valentinyi. 2014. "Growth and Structural Transformation." In *Handbook of Economic Growth*. Vol. 2 of *Handbook of Economic Growth*, Chapter 6, 855–941. Elsevier.

Hill, Martha. 1985. "Patterns of Time Use." In *Time, Goods, and Well-Being.* , ed. Thomas Juster & Frank Stafford. University of Michigan.

Jaimovich, Nir, & Henry Siu. 2012. "The Trend is the Cycle: Job Polarization and Jobless Recoveries." National Bureau of Economic Research Working Paper 18334.

Juhn, Chinhui, & Simon Potter. 2006. "Changes in Labor Force Participation in the United States." *Journal of Economic Perspectives*, 20(3): 27–46.

Kambourov, Gueorgui, & Iourii Manovskii. 2009. "Occupational Specificity Of Human Capital." *International Economic Review*, 50(1): 63–115.

Katz, Lawrence, & Kevin Murphy. 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *The Quarterly Journal of Economics*, 107(1): 35–78.

Kongsamut, Piyabha, Sergio Rebelo, & Danyang Xie. 2001. "Beyond Balanced Growth." *The Review of Economic Studies*, 68(4): 869–882.

Kuznets, Simon. 1941. *National Income and Its Composition, 1919-1938.* National Bureau of Economic Research, Inc. with Lillian Epstein & Elizabeth Jenks.

Mandelman, Federico. 2016. "Labor Market Polarization and International Macroeconomic Dynamics." *Journal of Monetary Economics*, 79: 1 – 16.

Ngai, Rachel, & Barbara Petrongolo. 2014. "Gender Gaps and the Rise of the Service Economy." Centre for Macroeconomics (CFM) Discussion Papers 1404.

Ngai, Rachel, & Christopher Pissarides. 2007. "Structural Change in a Multisector Model of Growth." *American Economic Review*, 97(1): 429–443.

Ngai, Rachel, & Christopher Pissarides. 2008. "Trends in Hours and Economic Growth." *Review of Economic Dynamics*, 11(2): 239–256.

Prescott, Edward. 2004. "Why do Americans Work so Much More than Europeans?" *Quarterly Review*, 28(1): 2–13.

Rogerson, Richard. 2009. "Market Work, Home Work, and Taxes: A Cross-Country Analysis." *Review of International Economics*, 17(3): 588–601.

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, & Matthew Sobek. 2015. "Integrated Public Use Microdata Series: Version 6.0." Minneapolis: University of Minnesota [Machine-readable database].

Tüzemen, Didem, & Jonathan Willis. 2013. "The Vanishing Middle: Job Polarization and Workers' Response to the Decline in Middle-Skill Jobs." *Economic Review*, 98(1): 5–32.

Appendix A Data Sources

The data I use covers the 1968-2018 period, and comes from the employment data in the Annual Social and Economic (ASEC) supplement to the Current Population Survey. I accessed these databases from the IPUMS-CPS project, an integrated set of data from the Current Population Surveys that goes through a convenient harmonization process. I consider the population aged between 25 and 65 years, and use their labels to determine industry, occupation (for the employed), and labor force status.

I construct the industry categories by grouping into goods and services. The goods industry includes both the manufacturing and agriculture sectors, which encompass manufacturing, construction, mining, agriculture, forestry, and fishery industrial classifications. The industrial classifications of the services industry are transportation, communications, public utilities, wholesale trade, retail trade, finance, insurance, real estate, business and repair services, personal services, entertainment and recreation services, professional and related services, and public administration.

To construct the occupation categories, I follow Cortés et al. (2014). They classify occupations based on two criteria: whether the tasks they involve are primarily manual or cognitive, and whether these are of a routine nature or not. Jaimovich & Siu (2012) note that there is a ranking in terms of wages: non-routine cognitive earn the highest while non-routine manual the lowest. In terms of tasks, non-routine cognitive tend to be high-skilled, while non-routine manual tend to be low-skilled. Routine manual and routine cognitive tend to be middle-skilled, so I group them together, in a similar fashion to Cortés (2016). I end up with three occupation groups then: non-routine manual, routine, and non-routine cognitive. Due to their association with the skills required, in the rest of the article I refer to these as manual, routine, and abstract occupations.

With every decennial Census, the occupation classifications are revised. These imply discrete jumps in their structure; even with the coarse grouping I use the changes are visible. The biggest changes were made with the 1983 and 2003 Censuses, which explain some of the shifts in the figures presented later. Both the harmonization processes from the IPUMS project and the analyses in Cortés et al. (2014) are careful enough to try and minimize these effects. In terms of the longer time trends, these reclassifications do not alter overall patters, and do not represent a significant concern.

Appendix B Detailed Tables

Table 8: Occupational Job Polarization

	1968	2018
Manual	11.2	14.2
Routine	59.3	42.9
Abstract	29.5	42.9

These percentages refer to each occupation's share in employment.
Source: author's calculations using CPS.

Table 9: Job Polarization and Non-employment

	1968	2018
Manual	7.6	11.1
Routine	40.0	33.7
Abstract	19.9	33.6
Non-employment	32.6	21.6

These percentages refer to each category's share in the total population.
Source: author's calculations using CPS.

Table 10: Occupation Shares within Industries

	Goods		Services	
	1968	2018	1968	2018
Manual	1.2	1.3	18.1	17.6
Routine	76.7	68.5	47.4	36.2
Abstract	22.1	30.2	34.5	46.2

These percentages refer to the share of each occupation the industry's labor demand.
Source: author's calculations using CPS.

Table 11: Industry Shares

	1968	2018
Goods	40.6	20.9
Services	59.4	79.1

These percentages refer to the share of each industry in the labor force.
Source: author's calculations using CPS.

Appendix C Model

In this section, I present a static model of labor allocation between occupations to study the patterns shown earlier. It is a variation of Duernecker & Herrendorf (2016) and Ngai & Petrongolo (2014): a model of structural transformation, that features occupational choices within the firms, and allows for labor non-participation by including a home production sector.

The agents in this model choose between market and non-market work. In the market, firms decide how to allocate their labor into the three market occupations: manual, routine or abstract. In non-market work, agents devote time exclusively to home production. The driving force is *occupation specific* technical progress, and the difference in their growth rates induces the three main results: polarization, structural transformation, and changes in the labor force participation.

C.1 Environment

This is a discrete-time model where time runs forever. On the production side, I follow Ngai & Petrongolo (2014) and study three productive sectors: goods, market services, and home services. To distinguish between the jobs agents are working in, and the industries where these take place, I denote by lowercase j the occupation, and by uppercase I the industry. Then, $j \in \{h, m, r, a\}$, meaning these jobs can be in home production, manual occupations, routine occupations, and abstract occupations. Similarly, $I \in \{G, M, H\}$ denotes the production of goods, of market services, and of home services.

Home services are produced with a linear technology on home labor:

$$Y_{Ht} = A_{ht} N_{Hht} \tag{17}$$

where A_{ht} denotes the efficiency of home production, and N_{Hht} denotes the amount of labor used in home production.

The firms in goods and market services produce with a technology that requires the three types of market occupations: manual, routine, and abstract. These are combined according to a CES aggregator:

$$Y_{It} = \left[\sum_{j \in \{m,r,a\}} \alpha_{Ij}^{\frac{1}{\sigma}} (A_{jt} N_{Ijt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{18}$$

where $I \in \{G, M\}$ indicates the industry and $j \in \{m, r, a\}$ the occupation. In this setting, N_{Ijt} denotes the input that industry I uses of occupation j, and A_{jt} is the labor efficiency in occupation j. The elasticity of substitution between the labor inputs is $\sigma > 0$, and $\alpha_{Ij} \in (0, 1)$ is the intensity of occupation j in sector I. This productive structure is similar to Duernecker & Herrendorf (2016), where labor efficiency is occupation-specific as opposed to industry specific, which is the standard assumption in the structural transformation literature.

For notational convenience, define the following:

$$N_{It} = N_{Imt} + N_{Irt} + N_{Iat} \qquad I \in \{G, M\} \tag{19}$$

$$N_{jt} = N_{Gjt} + N_{Mjt} \qquad j \in \{m, r, a\} \tag{20}$$

$$\begin{aligned} N_t &= N_{mt} + N_{rt} + N_{at} \\ &= N_{Gt} + N_{Mt} \end{aligned} \tag{21}$$

where N_{It} is the total amount of labor in industry $I \in \{G, M\}$, N_{jt} is the total amount of labor in occupation $j \in \{m, r, a\}$, and N_t is total market labor, which is clearly equal to the sum of labor over market industries or occupations.

On the consumption side, there are identical households of measure one. These consume goods and a combination of home and market services. The utility level they yield is aggregated according to a nested CES specification:

$$U_t(C_{Gt}, C_{St}) = \left[\omega_G^{\frac{1}{\varepsilon}} (C_{Gt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_S^{\frac{1}{\varepsilon}} (C_{St})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{22}$$

where C_{Gt} and C_{St} denote the consumption of goods and compound services. Their relative weights are ω_G and ω_S, which add up to one, and individually are between zero and one. The elasticity of substitution between goods and compound services is $\varepsilon > 0$.

Compound services are also aggregated through a CES specification:

$$C_{St} = \left[\varphi_M^{\frac{1}{\eta}}(C_{Mt})^{\frac{\eta-1}{\eta}} + \varphi_H^{\frac{1}{\eta}}(C_{Ht})^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}} \tag{23}$$

where C_{Mt} and C_{Ht} denote market and home services. Their relative weights are φ_M and φ_H, which also add up to one, and individually are between zero and one. The elasticity of substitution between market and home services is $\eta > 0$.

All households are endowed with one unit of labor in each period. I denote by L_t total labor supply, and the remaining $1 - L_t$ is devoted to home production. Within the market, labor is perfectly mobile across occupations, and has no occupation or sector specificity to it.

The feasibility conditions for the consumption sectors are:

$$Y_{Gt} = C_{Gt} \tag{24}$$

$$Y_{Mt} = C_{Mt} \tag{25}$$

$$Y_{Ht} = C_{Ht} \tag{26}$$

These equations simply require that what is produced in the goods, market services and home services industries be consumed by the households.

Finally, the feasibility condition for the labor market requires that the households' labor supply be equal to the market demand:

$$\begin{aligned} L_t &= N_{mt} + N_{rt} + N_{at} \\ &= N_{Gt} + N_{Mt} \\ &= N_t \end{aligned} \tag{27}$$

C.2 Decentralized Market Structure

I assume competitive markets for labor, goods and market services, where all agents take the prices as given. Firms in the consumption and market services industries face the following profit-maximization problem:

$$\max_{N_{Imt},N_{Irt},N_{Iat}} p_{It}\left[\sum_{j\in\{m,r,a\}}\alpha_{Ij}^{\frac{1}{\sigma}}(A_{jt}N_{Ijt})^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}} - w_t(N_{Imt}+N_{Irt}+N_{Iat}) \tag{28}$$

where p_{It} is the market price of their output, and w_{jt} the market wages. First order conditions imply:

$$\frac{N_{Imt}}{N_{Irt}} = \frac{\alpha_{Im}}{\alpha_{Ir}}\left(\frac{A_{rt}}{A_{mt}}\right)^{1-\sigma} \tag{29}$$

$$\frac{N_{Iat}}{N_{Imt}} = \frac{\alpha_{Ia}}{\alpha_{Im}}\left(\frac{A_{mt}}{A_{at}}\right)^{1-\sigma} \tag{30}$$

Equations (29) and (30) describe the relative labor allocations between manual and routine occupations, and abstract and manual occupations, each for industry $I \in \{G, M\}$.

Taking prices as given, the household's utility-maximization problem at time t is:

$$\max_{L_t,C_{Gt},C_{Mt}} \left[\omega_G^{\frac{1}{\varepsilon}}(C_{Gt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_S^{\frac{1}{\varepsilon}}(C_{St})^{\frac{\varepsilon-1}{\varepsilon}}\right]^{\frac{\varepsilon}{\varepsilon-1}} \tag{31}$$

subject to:

$$p_{Gt}C_{Gt} + p_{Mt}C_{Mt} = w_tL_t$$

$$C_{St} = \left[\varphi_M^{\frac{1}{\eta}}(C_{Mt})^{\frac{\eta-1}{\eta}} + \varphi_H^{\frac{1}{\eta}}(C_{Ht})^{\frac{\eta-1}{\eta}}\right]^{\frac{\eta}{\eta-1}}$$

$$C_{Ht} = A_{ht}(1-L_t)$$

Households then, maximize their utility subject to their budget constraint in market products, and their technology constraint in household production. This problem can be solved in two steps: the first is to find the optimal allocation between home and market services, and the second one is for the optimal allocation between goods and compound services.

Home services are not traded in the market, meaning there is no market price attached to them. Its opportunity cost, however, is well defined since its alternatives have market prices attached to them. I denote by p_{Ht} this implicit price. The first order conditions to maximize C_{St} imply:

$$\frac{p_{Mt}C_{Mt}}{p_{Ht}C_{Ht}} = \frac{\varphi_M}{\varphi_H}\left(\frac{p_{Mt}}{p_{Ht}}\right)^{1-\eta} \tag{32}$$

Define the following price index:

$$p_{St} = \left[\varphi_H(p_{Ht})^{1-\eta} + \varphi_M(p_{Mt})^{1-\eta}\right]^{\frac{1}{1-\eta}} \tag{33}$$

This price index can be interpreted as the unit price of the optimal services basket, which is relevant for the decision between the consumption of goods and the composite services basket. For this decision, first order conditions imply:

$$\frac{p_{St}C_{St}}{p_{Gt}C_{Gt}} = \frac{\omega_S}{\omega_G}\left(\frac{p_{St}}{p_{Gt}}\right)^{1-\varepsilon} \tag{34}$$

Appendix D Estimation Procedure

In this section, I explain in further detail how to match the data to the model's parameters. Recall the assumption of constant growth rates in productivity; because of that I only need to look at the initial and final years. These are denoted by $t = 0$ and $t = T$.

To get the market occupation intensities in production (α_{Ij}), I use equations (29) (30). For the initial year, these imply:

$$\alpha_{Ir} = \alpha_{Im}\frac{N_{Ir0}}{N_{Im0}} \qquad \qquad \alpha_{Ia} = \alpha_{Im}\frac{N_{Ia0}}{N_{Im0}} \tag{35}$$

These two yield the intensities, since all three add up to one.

To get the relative weights in the consumption of market and home services (φ_H and φ_M), and in the consumption of goods and services (ω_G and ω_S), I use equations (13) and (15) in a similar fashion:

$$\varphi_H = \varphi_M\frac{N_{H0}}{N_{M0}} \qquad \qquad \omega_G = \omega_S\varphi_M\frac{N_{G0}}{N_{M0}} \tag{36}$$

To get the elasticity of substitution in the production function, I first rewrite the average labor productivity (10):

$$\begin{aligned}\tilde{A}_{It} &= A_{rt}\left\{\alpha_{Ir}\left[1+\frac{\alpha_{Im}}{\alpha_{Ir}}\left(\frac{A_{rt}}{A_{mt}}\right)^{1-\sigma}+\frac{\alpha_{Ia}}{\alpha_{Ir}}\left(\frac{A_{rt}}{A_{at}}\right)^{1-\sigma}\right]\right\}^{\frac{-1}{1-\sigma}} \\ &= A_{rt}\left\{\alpha_{Ir}\left[1+\frac{N_{Imt}}{N_{Irt}}+\frac{N_{Iat}}{N_{Irt}}\right]\right\}^{\frac{-1}{1-\sigma}} \end{aligned} \tag{37}$$

Then, I use equation (15) in the final period, substituting in equations (13) and (37):

$$\frac{N_{MT}}{N_{GT}} = \frac{\omega_S}{\omega_G}\left[\frac{\alpha_{Mr}}{\alpha_{Gr}}\frac{\left(1+\frac{N_{MmT}}{N_{MrT}}+\frac{N_{MaT}}{N_{MrT}}\right)}{\left(1+\frac{N_{GmT}}{N_{GrT}}+\frac{N_{GaT}}{N_{GrT}}\right)}\right]^{\frac{1-\varepsilon}{1-\sigma}}\left(\varphi_M\frac{N_{ST}}{N_{MT}}\right)^{\frac{1-\varepsilon}{1-\eta}} \tag{38}$$

Applying logarithms and rearranging leads to my estimate of σ.

To get relative occupational productivities, I use the occupation shares in period T, and equations (29) and (30) for the market services industry. Rearranging, these give:

$$\frac{A_{mT}}{A_{rT}} = \left(\frac{\alpha_{Mm}}{\alpha_{Mr}}\frac{N_{MrT}}{N_{MmT}}\right)^{\frac{1}{1-\sigma}} \qquad\qquad \frac{A_{aT}}{A_{rT}} = \left(\frac{\alpha_{Ma}}{\alpha_{Mr}}\frac{N_{MrT}}{N_{MaT}}\right)^{\frac{1}{1-\sigma}} \tag{39}$$

To get the productivity levels in market occupations, I use data from the Bureau of Economic Analysis. In particular, I take the real gross domestic product per capita time series (A939RX0Q048SBEA) to establish that this has grown by a factor of 2.23 between 1968 and 2018. To reproduce this growth pattern, I match this factor with market production evaluated at year 0's prices. Substituting (37) into the production function (18):

$$\begin{aligned} 2.23 &= \frac{p_{M0}Y_{MT}+p_{G0}Y_{GT}}{p_{M0}Y_{M0}+p_{G0}Y_{G0}} \\ &= \frac{A_{rT}}{A_{r0}}\frac{N_{MT}\left[\alpha_{Mr}\left(1+\frac{N_{MmT}}{N_{MrT}}+\frac{N_{MaT}}{N_{MrT}}\right)\right]^{\frac{-1}{1-\sigma}}+N_{GT}\left[\alpha_{Gr}\left(1+\frac{N_{GmT}}{N_{GrT}}+\frac{N_{GaT}}{N_{GrT}}\right)\right]^{\frac{-1}{1-\sigma}}}{N_{M0}\left[\alpha_{Mr}\left(1+\frac{N_{Mm0}}{N_{Mr0}}+\frac{N_{Ma0}}{N_{Mr0}}\right)\right]^{\frac{-1}{1-\sigma}}+N_{G0}\left[\alpha_{Gr}\left(1+\frac{N_{Gm0}}{N_{Gr0}}+\frac{N_{Ga0}}{N_{Gr0}}\right)\right]^{\frac{-1}{1-\sigma}}} \end{aligned} \tag{40}$$

This yields the growth factor of productivity in routine occupations. With this I reconstruct the other productivity levels for market occupations.

Finally, from equation (13) in the final year, I get the productivity level in home production:

$$\frac{A_{rT}}{A_{hT}} = \left[\frac{\varphi_M}{\varphi_H}\frac{N_{HT}}{N_{MT}}\right]^{\frac{1}{1-\eta}} \left[\alpha_{Mr}\left(1 + \frac{N_{MmT}}{N_{MrT}} + \frac{N_{MaT}}{N_{MrT}}\right)\right]^{\frac{1}{1-\sigma}} \tag{41}$$